Analysis of DITL root data and comparison with a full-resolver's data

Kazunori Fujiwara

JPRS and University of Tsukuba

<fujiwara@jprs.co.jp>

DNS-OARC 2014 Spring Workshop

Updates from my previous analysis

- Analyzed DITL data from 2006 to 2010
  - Added some new graphs
- Compared with full-resolver data in 2012
  - A full-resolver sent 100,000 queries to root within 48 hours
  - It may be a typical full-resolver
- Previous analysis was reported at DNS-OARC 2013 Fall Workshop

Datasets and analysis method

DNS-OARC Root Datasets (1)

- "A Day in the Life of the Internet" (DITL) is a large-scale data collection project undertaken by CAIDA and DNS-OARC every year since 2006.
  - https://www.dns-oarc.net/oarc/data/ditl
  - 50 hours packet capture at root DNS servers and other DNS servers (48 hours are used by this analysis)
  - Source IP addresses of i.root-servers.net data are anonymized

DNS-OARC Root Datasets (2)

Year	Start (UTC)	End	List of root servers
2006	Jan 10 0000	Jan 12 0100	c,e,f,k (4/13)
2007	Jan 09 0000	Jan 11 0000	c,f,k,m (4/13)
2008	Mar 18 0000	Mar 20 0000	a,c,e,f,h,k,l,m (8/13)
2009	Mar 30 0000	Apr 02 0000	a,c,e,f,h,k,l,m (8/13), 72 hours
2010	Apr 14 0000	Apr 16 0000	a,b,c,d,e,f,g,h,i,j,k,l,m (12/13)
2011	Apr 12 1200	Apr 14 1200	a,c,d,e,f,h,j,k,l,m (10/13)
2012	Apr 17 1200	Apr 19 1200	a,c,e,f,h,j,k,l,m (9/13)
2013	May28 1200	May30 1200	a,c,d,e,f,h,j,k,l,m (10/13)

Analysis method of Root data

- Newly developed C program reads pcap files
- It counts number of some kinds of queries per IP address
  - All queries, RD=0 queries, EDNS0 queries,
  - DO set queries, name error queries,
  - "." DNSKEY queries (RD=0), “.” NS queries,
  - "." Queries, UDP checksum off queries
  - Port number bitmaps (to analyze source port randomization trends)
  - TLD bitmaps

University of Tsukuba Dataset

- Associate researchers and the author collected packet captures at one of full-resolvers at University of Tsukuba
- Around January 2011 to August 2012
- A data exist at the same timing as DITL-2012
  - Apr 17 12:00 to Apr 19 12:00 UTC
  - 72,355,778 DNS packets captured (418 pps)
  - 28,815,955 stub queries observed (166 qps)
  - 8429 unique query source addresses

Analysis method of full-resolver data

- Newly developed perl program
  - reads pcap files
  - classifies packets into
    - Stub queries/responses
    - Full-resolver to authoritative queries/responses
  - classifies destination addresses into
    - Root, TLDs, other authoritative servers
  - parses each section
  - parses referrals (NS+glue)
  - counts characteristics

Results

Number of IP addrs seen at root 48h

Year	2011		2012		2013	
Data from	10 root		9 root		10 root	
Total	7,591,031	100%	8,989,786	100%	8,547,065	100%
RD0	5,846,612	77.0%	5,859,493	65.2%	6,081,035	71.1%
EDNS0	2,340,543	30.8%	2,906,287	32.3%	3,572,804	41.8%
DO=1	2,018,839	26.6%	2,621,660	29.2%	3,283,728	38.4%
Update	105,131	1.4%	138,778	1.5%	228,633	2.7%
Update Only	71,972	0.9%	99,902	1.1%	179,874	2.1%
Non-existent TLD	2,606,340	34.3%	2,641,072	29.4%	2,619,836	30.7%
Existing TLD	7,361,794	97.0%	8,697,606	96.7%	8,142,126	95.3%
. NS	1,940,015	25.6%	1,871,995	20.8%	2,082,649	24.4%
. Only	26,877	0.4%	36,920	0.4%	105,784	1.2%
. DNSKEY (RD0)	14,092	0.2%	43,782	0.5%	269,390	3.2%
. DNSKEY . Only	571	0.0%	2,828	0.0%	64,612	0.8%

New graphs

- Horizontal axis: years from 2006 to 2013
- Graphs
  - Number of queries
  - Number of IP addresses
  - Ratio of IP addresses/queries
  - Port randomization status
- Data
  - Total, Update, EDNS0, DO
  - Repeat_port … an IP address sent from same port
    - one case of using static source port number
  - UdpSumOff
  - root: query name “.”
  - “.” DNSKEY

Number of queries

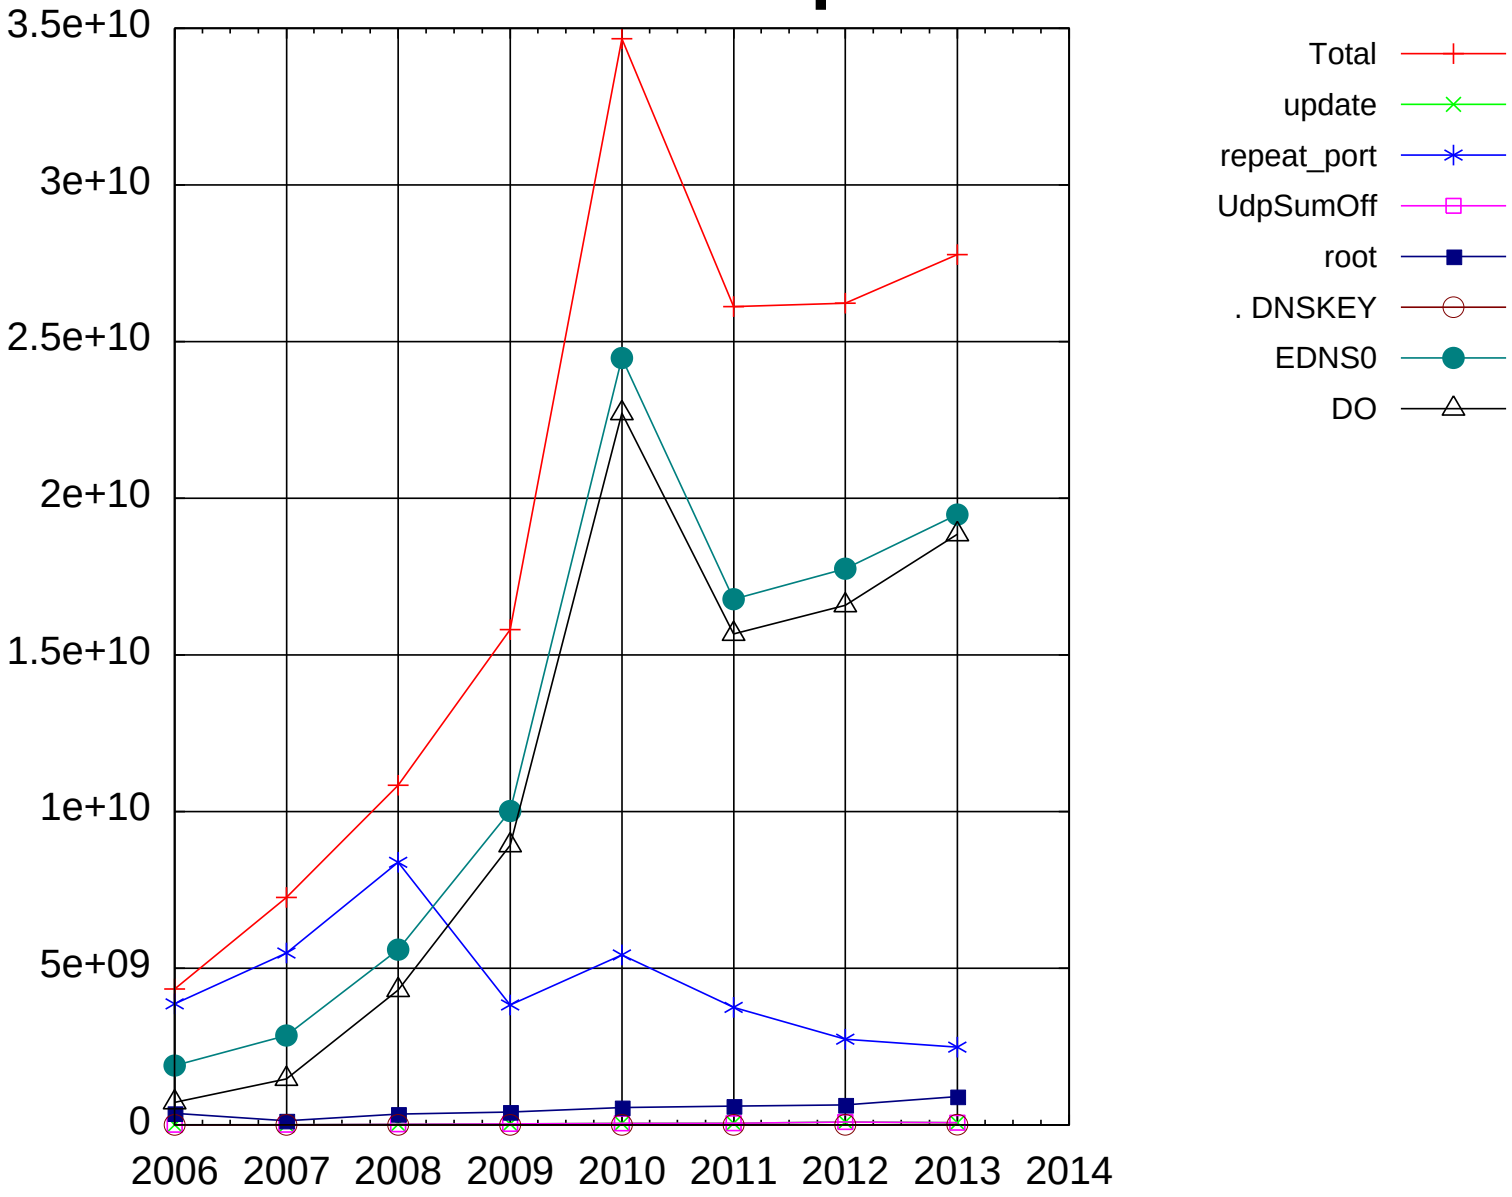

Number of IP addresses

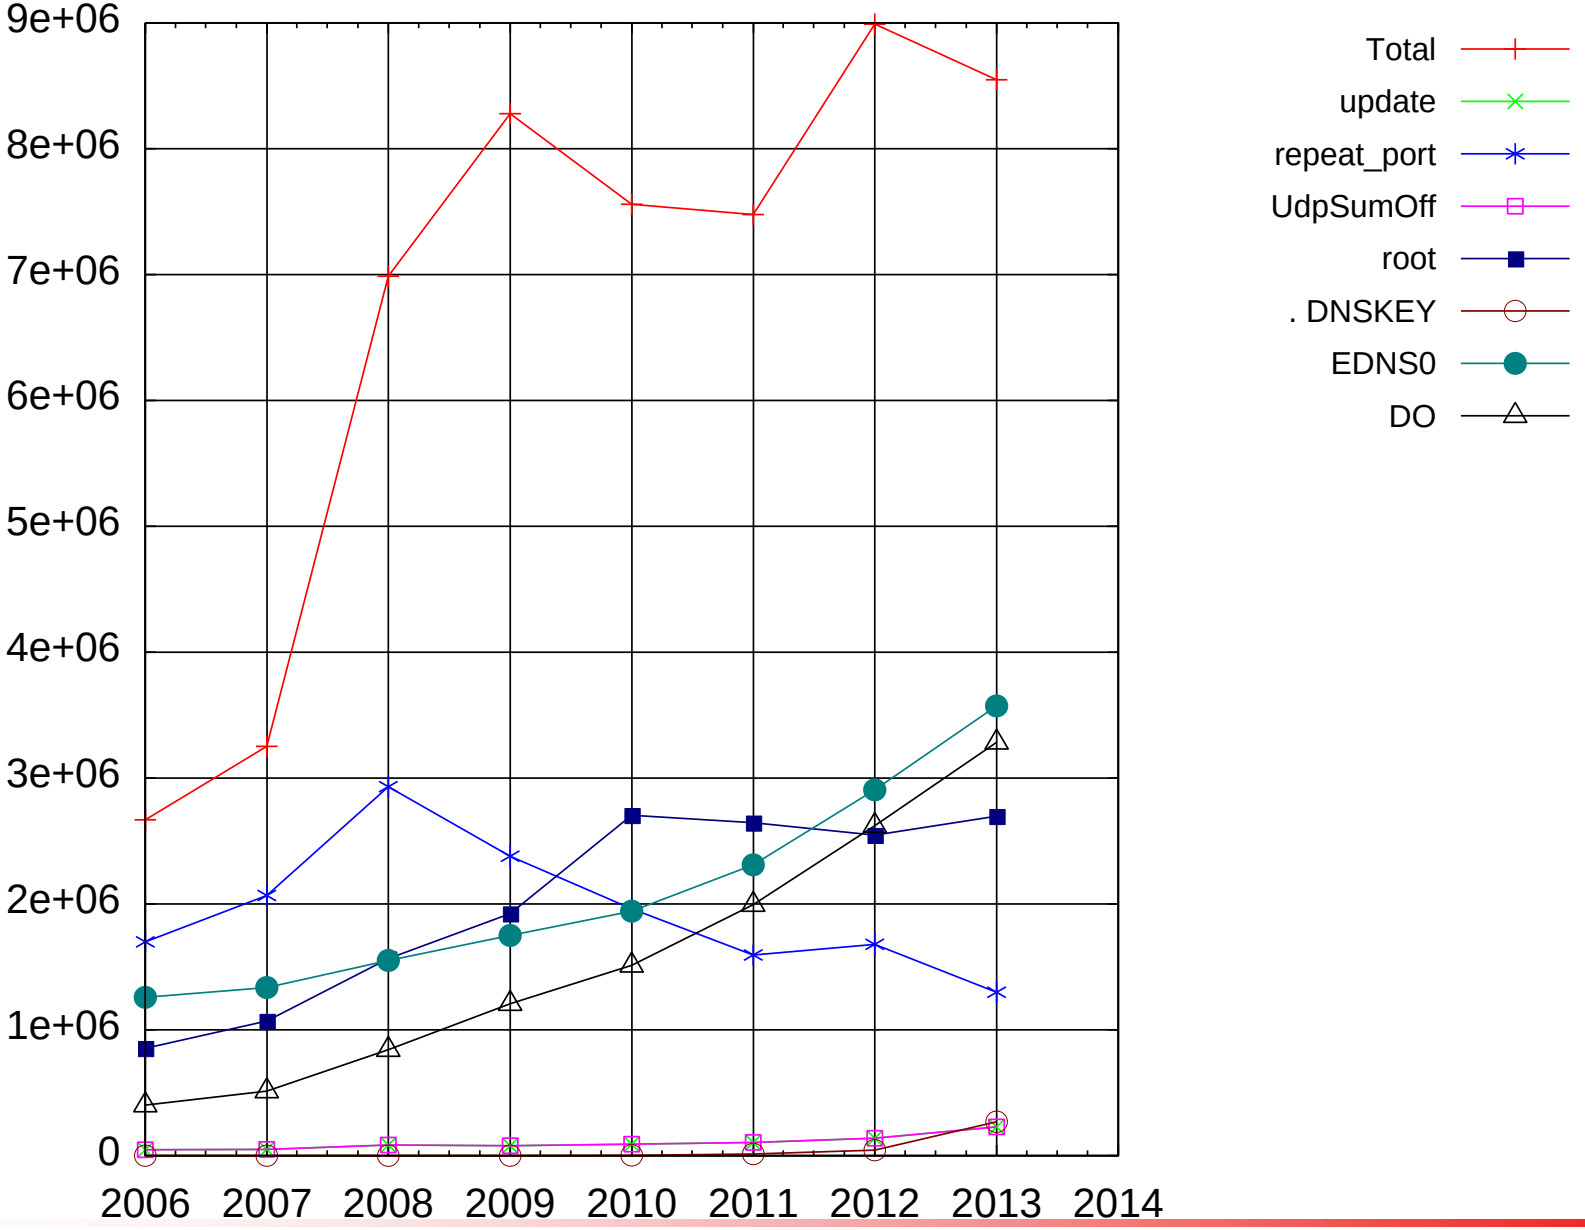

Ratio of addresses/queries

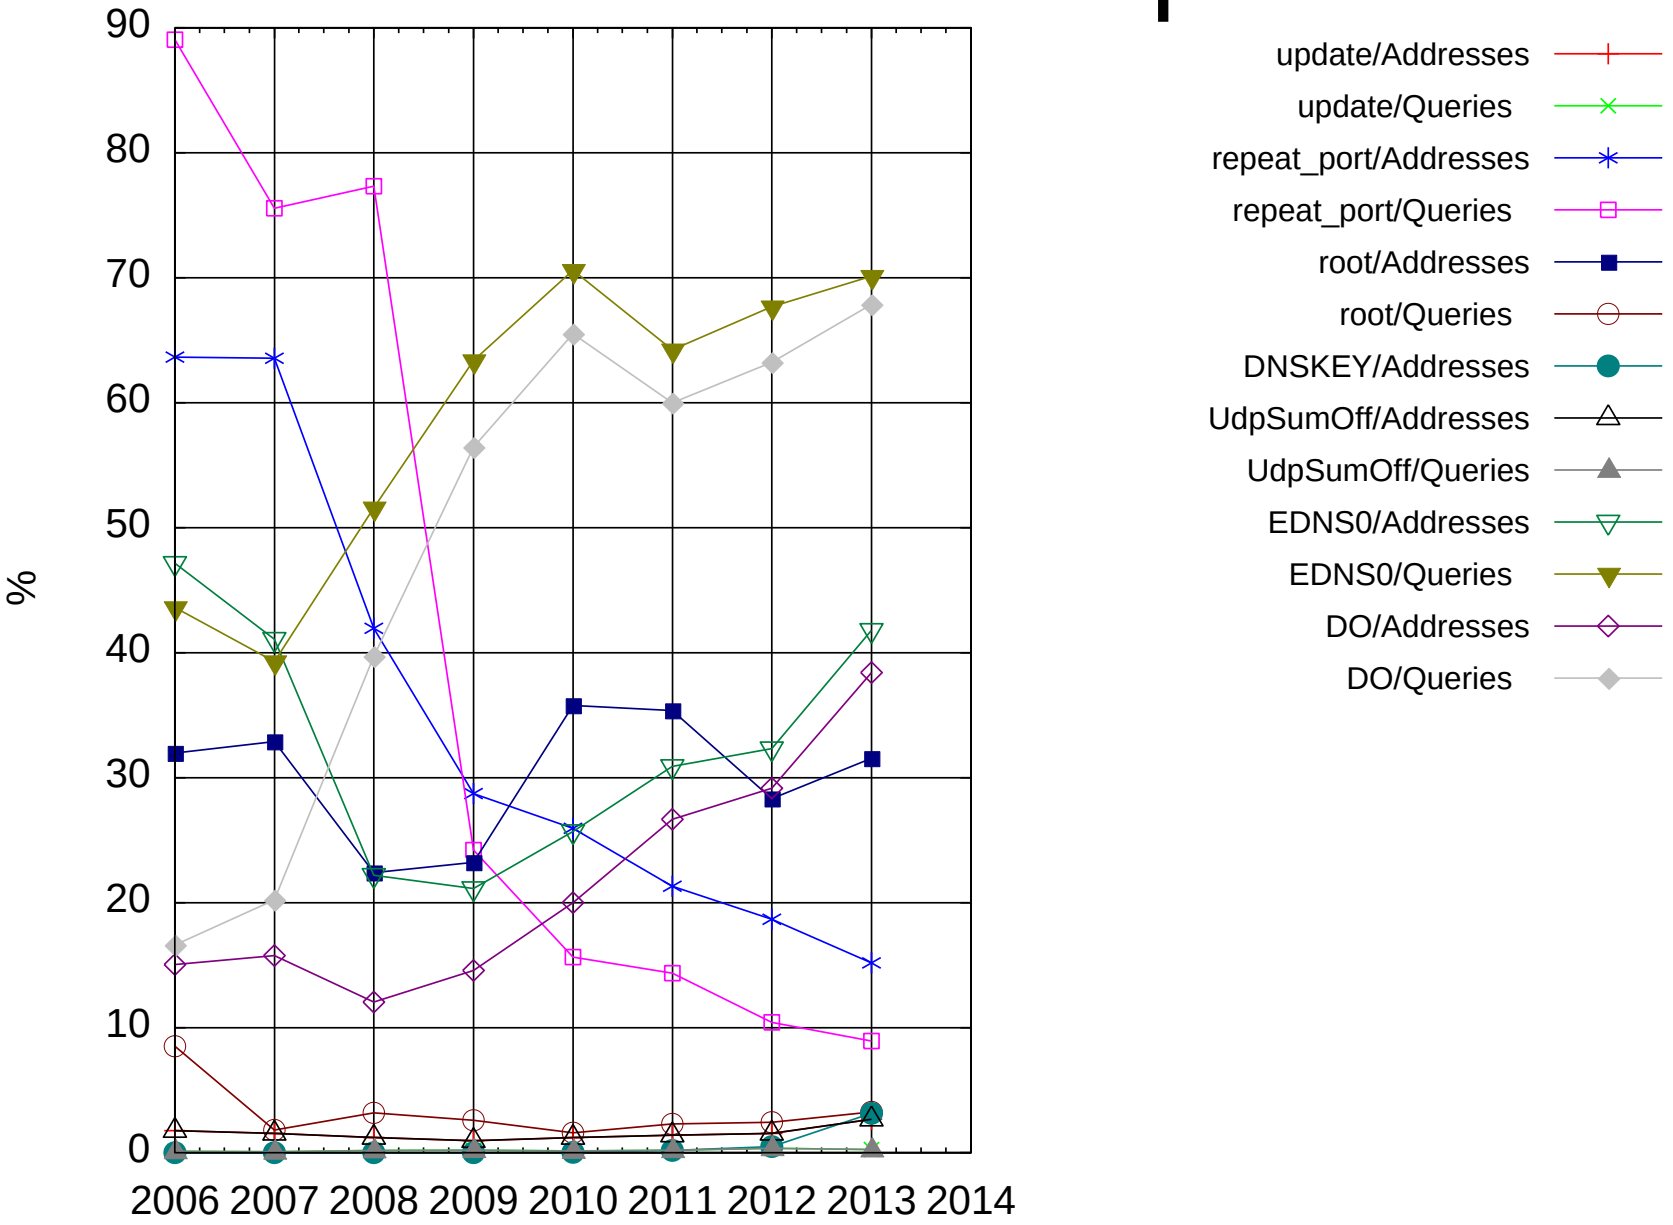

Status of port randomization

- High8 .. High 8 bits of port number
- Low8 .. Low 8 bits of port number
- OrderChange .. Number of changes of port numbers increase and decrease

- Unknown: queries from an IP address < 10
- Static: use of High8 < 4 and use of Low8 < 4
- Limited: use of High8 < 4 or use of Low8 < 4 or OrderChange < 4 (except Static)
- Random: others (port randomization enabled?)

- This classification is under concern

Source port randomization trends

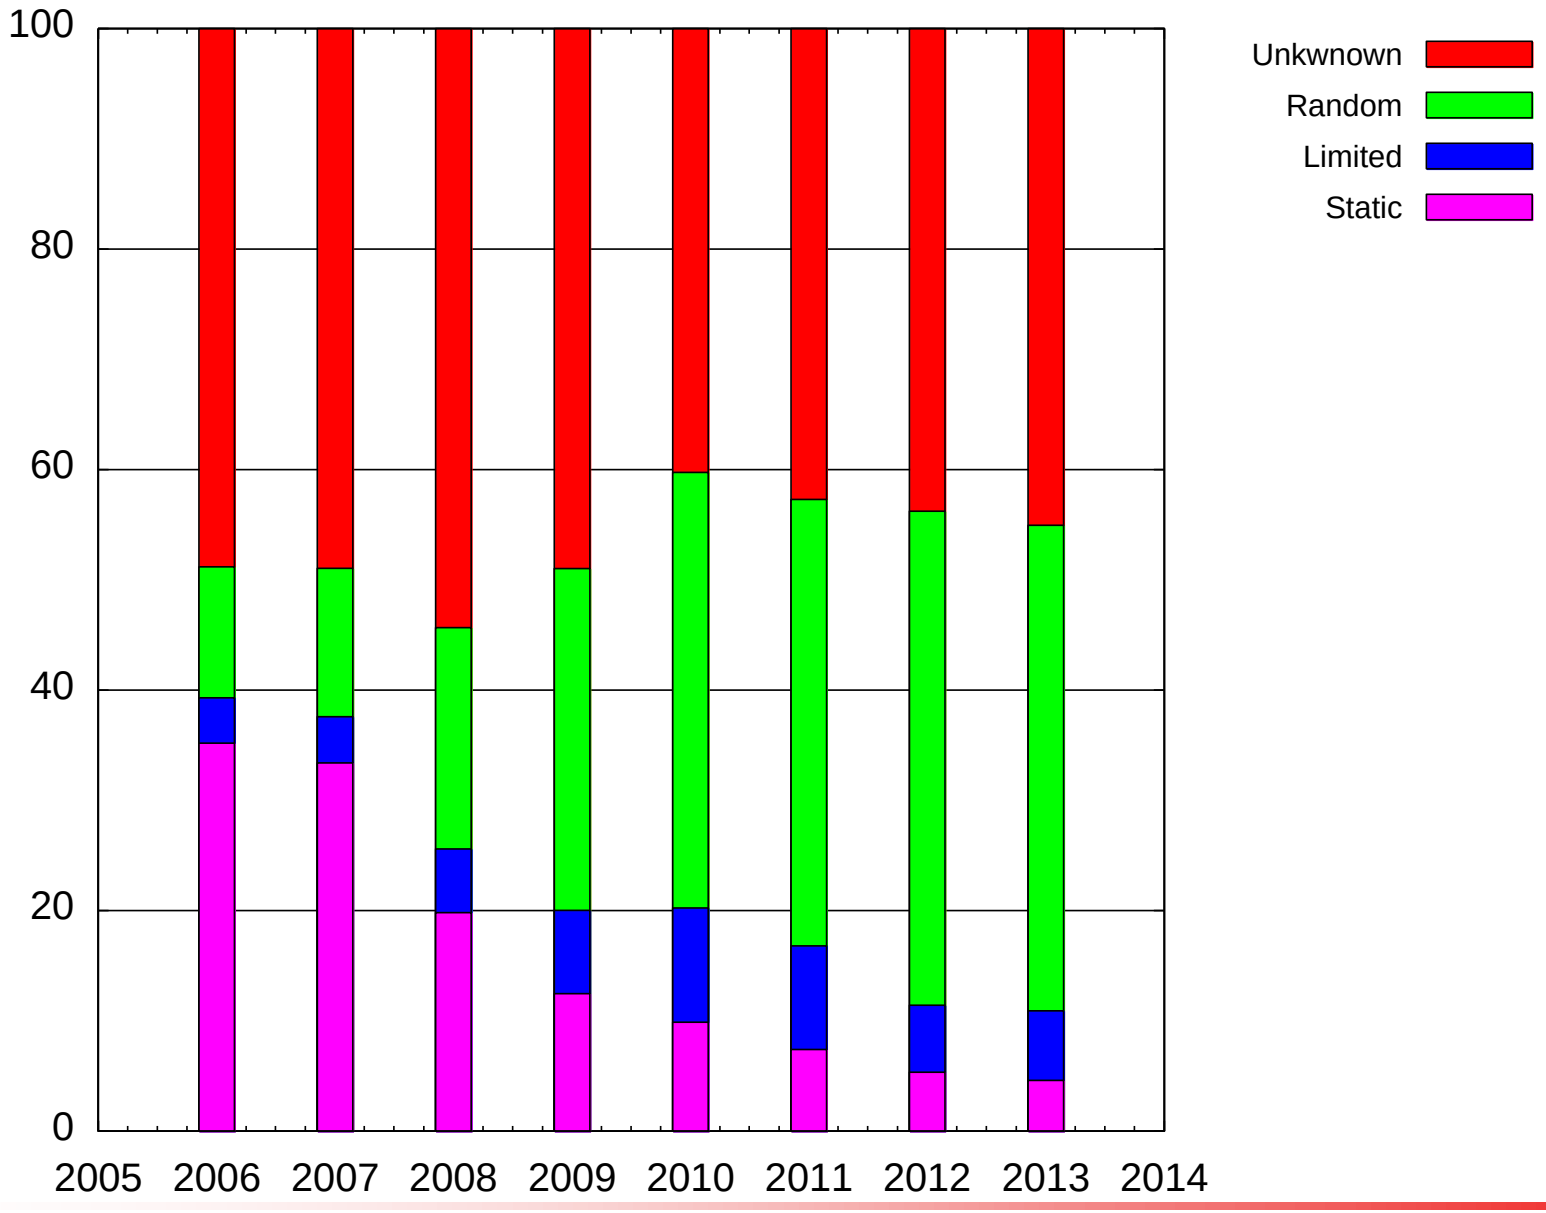

Number of queries send from each address, at root, 48 hours

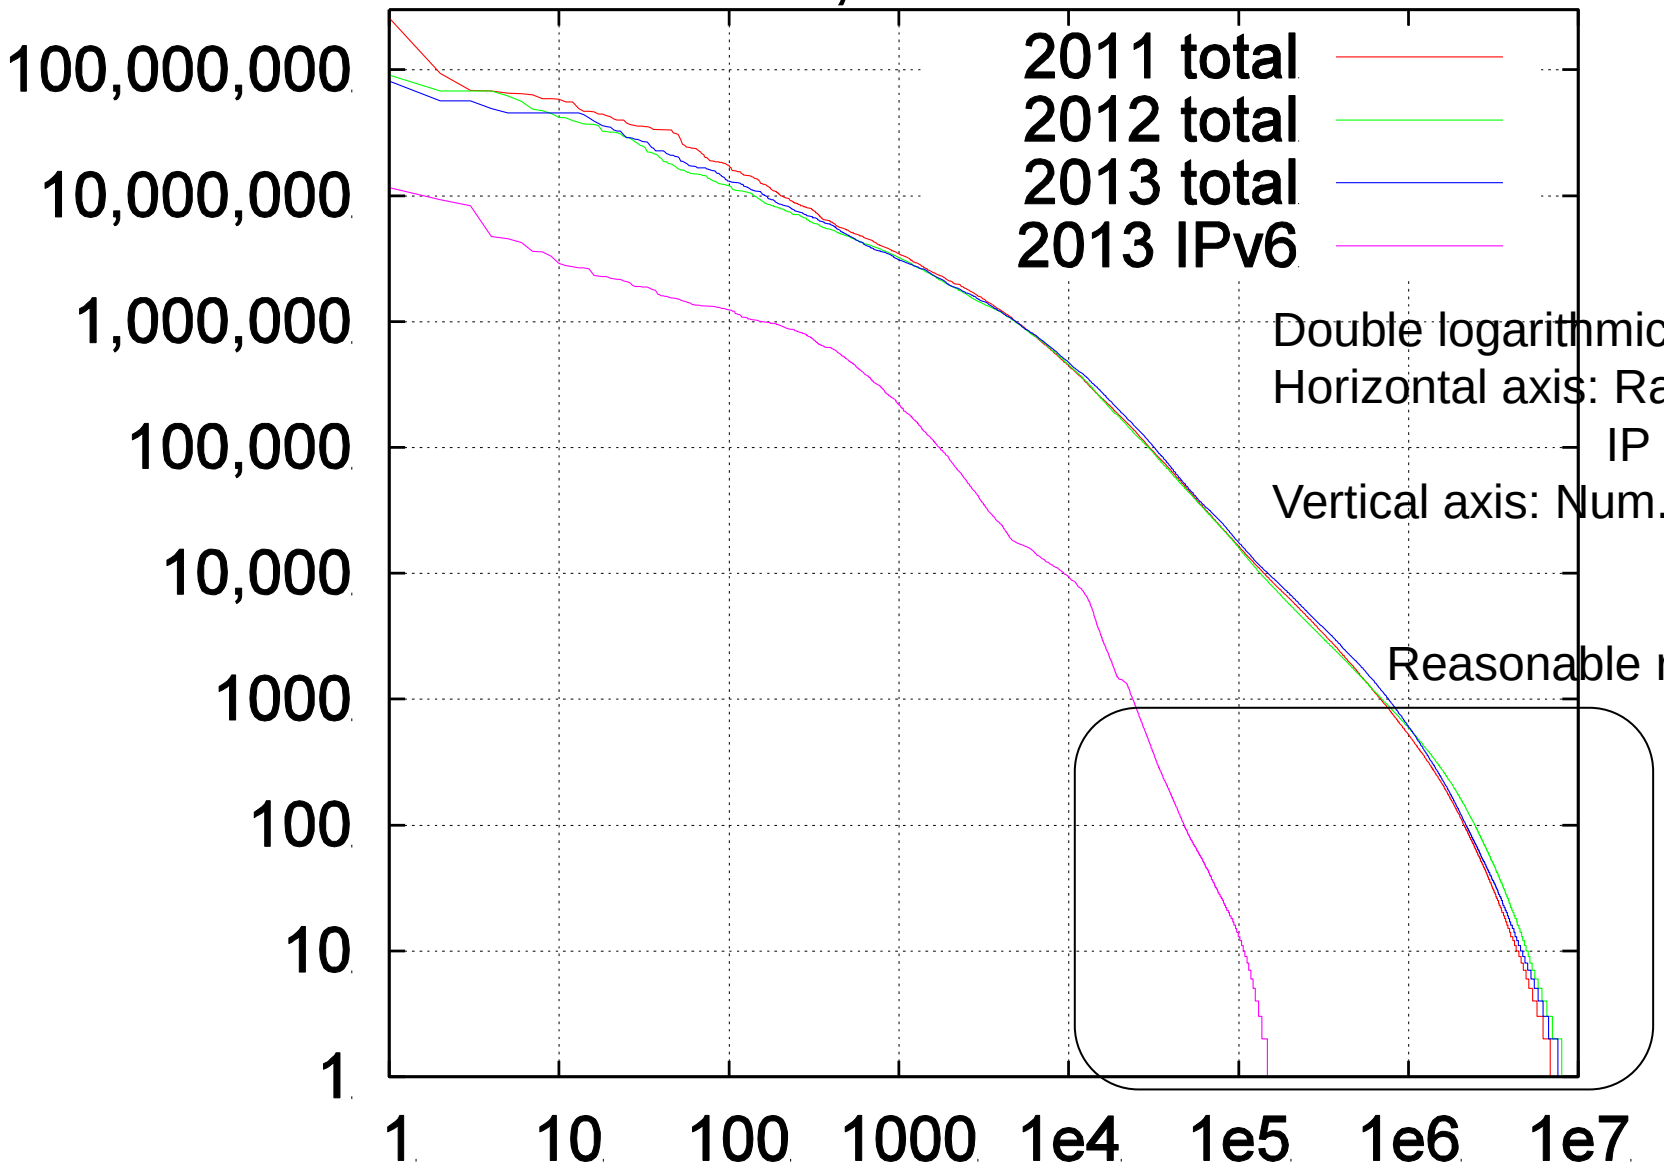

Double logarithmic chart
Horizontal axis: Ranking of IP addresses
Vertical axis: Num. of queries

Reasonable resolvers?

Number of queries from each address

- Without TLD typos,
- There were 318 TLDs and their NS TTLs were 172800 and DS TTLs were 86400 at May 2013
  - They should be cached within 1 or 2 days
- If resolvers work well, they should send only 2 * 318 + priming + root dnskey queries at most.
- However, there are 500,000 IP addresses which send over 1000 queries within 48 hours. Why ?
  - They send both existing names and non-existent names

Some known IP addresses

	Total queries to root	Non-existing TLD	"." queries	Existing name queries
My VPS (IPv4)	177	1	14	162
My VPS (IPv6)	182	2	12	168
My home server (IPv4)	1124	34	54	1036
One of JPRS address	21990	400	3	21587
One of full resolvers at University of Tsukuba	109215 Too many ?	12200	1298 Too Many?	95717 Too many?

Number of queries send from each address, at root, 48 hours

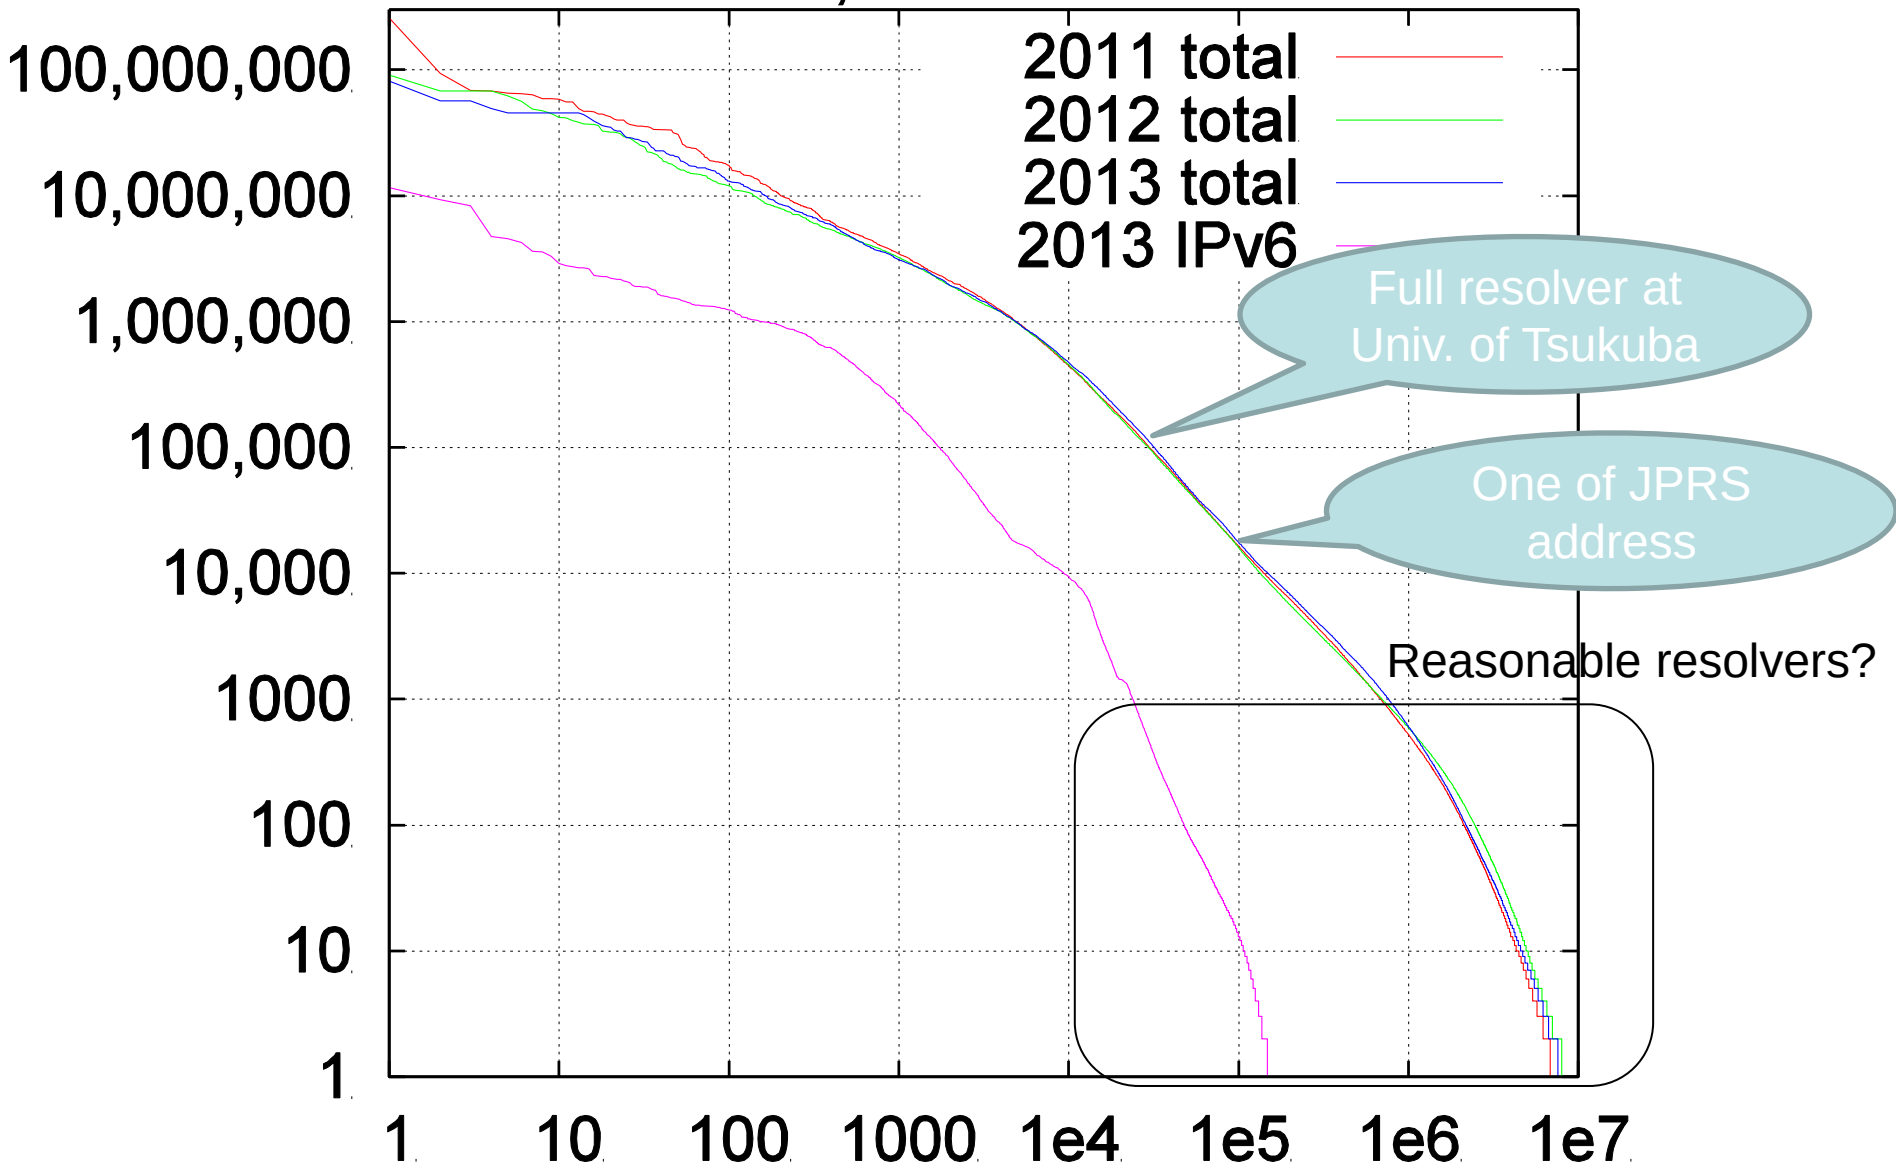

Status of the full-resolver at the university

- Software: latest version of BIND 9.6 (April 2012)
- Configuration
  - Recursion only
  - Without DNSSEC validation
  - Without special configurations

- There are packet captures of the full-resolver at the same timing of DITL-2012

Analysis of full-resolver packets

- 72,355,778 packets captured (418 pps)
- 28,815,955 stub queries (166 qps)
  - 1,026,487 non-existing TLD queries
- 8429 unique query source addresses
- 7,499,961 authoritative queries
- 7,329,795 authoritative answers
  - 118,360 root answers
    - 105,781 (89.4%) RCODE 0 Too many
    - 12,579 (10.6%) RCODE 3 Reasonable
  - 687,365 TLD answers

Observations from packet capture (1)

- At out-of-bailiwick delegations, a modern full- resolver will start resolving all DNS server names A and AAAA simultaneously
  - If the cache is empty, it will send twice as many queries as number of NS RRs to root
- TLD typos caused 12579 error responses from root
  - Most of them were different query names
  - Some came from small (negative) cacheable time (1 hour to 3 hours)

Observations from packet capture (2)

- The full-resolver at the university sent many DNS server name A/AAAA queries to root
  - Cacheable NS name queries are 103,024 (87%)
    - To mitigate attacks ?
    - Zone TTL is small ?
    - Why ?
- To understand this behavior more, I replayed client traffic to some full-resolvers to know that the behavior is true

Replay on some full-resolver software

- Input client traffic
  - 48 hours client query data at University of Tsukuba
  - Same timing, same qname/qtype
  - Original query source addresses are ignored
    - Sent from one IP address
- Tested full-resolvers
  - BIND 9.9.5 (with/without DNSSEC)
  - Unbound 1.4.22 (with/without DNSSEC, increased cache size, specify harden-referral)
    - large cache configuration
      - msg-cache-size: 1024m
      - rrset-cache-size: 1024m
      - infra-cache-numhosts: 500000

Results of replay

28,815,955 client queries

	To root	To root Name error	To root No error	To TLD	Auth
BIND 9.6 / Observed data	118,360	12,579	105,781	687,365	6,524,070
BIND 9.9.5	163,187	12,975	150,212	842,592	7,377,712
BIND 9.9.5 + DNSSEC	663,647	12,727	650,920	1,061,886	7,235,743
Unbound	99,923	25,914	74,309	3,916,313	16,048,069
Unbound + large cache	13,300	11,444	1,856	870,650	9,102,884
Unbound + large cache + harden referral path	20,897	11,234	9,663	2,328,026	11,152,425
Unbound + large cache + DNSSEC	12,662	11,140	1,522	1,423,789	9,112,902

Summary of replay

- It is preliminary result
  - Need more test and detailed analysis
  - Values change at every experiment
- Both BIND 9 and Unbound generated too many (around 100,000 / 48hours) root queries
  - 100,000 queries to root is real
- Queries to root increased 4 times when DNSSEC validation enabled on BIND 9
- Unbound with default configuration is not good for middle / large scale sites
  - Because Unbound with default configuration sent 74,309 positive queries to root
  - Unbound with large cache sent only 1,856 queries to root
- Using Unbound may decrease root queries. However queries to other servers may increase

Conclusion

- Port randomization is spreading gradually, however about 10% of IP addresses are still dangerous
- DITL data show that 30,000 IP addresses sent 100,000 or more queries to root DNS servers within 48 hours.
- A full-resolver at University of Tsukuba sent 118,360 queries within 48 hours to root.
- As a result of replay experiment, both BIND 9 and Unbound full-resolvers sent around 100,000 queries to root within 48hours
- BIND 9 full-resolvers may send many queries to root DNS servers
  - Unbound send smaller number of queries to root
  - Decreasing queries to root is important

Acknowledgements

- DNS-OARC as the data source of Root dataset
- Academic Computing & Communication Center offers Full-Resolver DNS servers for Campus network of University of Tsukuba